EXHIBIT 10

EXHIBIT
607
Jiandani
4/29/2016 C.A.R.

NETWORKWORLD

Arista, Blade win top spot in data center switch test

Grueling performance testing exposes weaknesses in Cisco, HP switches

By David Newman | Follow
Network World | Jan 18, 2010 12:00 AM PT

As data center managers consolidate and virtualize their servers, the next order of business becomes moving all that traffic. Enter top-of-rack data center switches that offer speed, scalability, redundancy, virtualization support and other features not available in garden-variety Ethernet switches.

How we tested the switches

Archive of Network World tests

Lies, Damned Lies and Latency

This test analyzes switches, each sporting at least 24 10Gigabit interfaces, from Arista Networks, Blade Network Technologies, Cisco, Dell, Extreme and HP. We compared these products 10 different ways and subjected them to three months of grueling performance tests.

While each offered some standout qualities, we're singling out Arista's DCS-7124S and Blade's G8124 as top picks. With the best combination of features and performance – especially in the areas of latency and jitter – both switches earn Clear Choice awards. The Dell and Extreme entries also fared well.

The biggest surprises came from Cisco and HP. Cisco's Nexus 5010 is the only switch tested with a complete story on data/storage convergence, and its lengthy features list includes some outstanding virtualization capabilities. But high latency, usability gremlins and multicast leakage all hampered the Nexus 5010 in this test. HP ProCurve's 6600-24XG was only a modest performer. Although it was tops in media access control (MAC) address scalability, it was the only switch not to offer line-rate throughput, either for unicast and multicast traffic.

To help compare switches, we asked vendors to complete an extensive features questionnaire (see results). While vendors responded to more than 100 questions, we'll focus here on the major differentiators.

First, though, it should be noted that we decided not to compare vendors on price because actual pricing varies so much from list price.

Comparing basic features

The Arista, Blade and Cisco switches accept both gigabit and 10G Ethernet transceivers, a useful feature as data centers migrate from the gigabit connections found in the majority of servers today to 10G links which are likely to appear in the next year or two, especially as these begin to be embedded on server motherboards.

Cisco's Nexus 5010 is the only switch tested that offers native Fibre Channel interfaces and full Fibre Channel over Ethernet (FCoE) support. Fibre Channel options include 2G, 4G and 8G versions. All other switches can forward FCoE traffic, but that's a bit like saying an Ethernet switch supports IPv6 even though it's unaware the frames it forwards contain IPv6 headers.

All switches support redundant power supplies, and all except Extreme's offer redundant fans. The Arista and Blade designs also offer redundant out-of-band management ports. Fans in the Arista, Cisco, Dell and HP switches are hot-swappable.

Airflow is a concern in data centers, and vendors take differing approaches to keep their switches cool. The Arista, Blade and HP switches can be configured to blow air from front to back or vice-versa, depending on whether the switch faces a hot or cold row. Cisco's switch also reverses the orthodox design so that interfaces are at the back of the switch, alongside the power supplies; Cisco says this minimizes clutter.

Although data center switches that connect virtual machines will operate mainly in layer-2 switching mode, many network designs involve at least some layer-3 routing. In this test, all devices except Cisco's Nexus 5010 can route packets between subnets, either statically or using open shortest path first (OSPF). The Blade, Dell and Extreme switches also support OSPFv3 for routing of IPv6 packets.

Data center switches support failover protocols on server and inter-switch links that not only boost uptime but also may eliminate the need for other redundancy protocols such as spanning tree.

For both server and inter-switch connectivity, all switches support the use of link aggregation to bond multiple physical connections into a single logical link.

Some vendors go beyond this with supplements or replacements for spanning tree on inter-switch links. Arista's multiple link aggregation (MLAG) technology allows active/active, multi-homed connections for server and inter-switch connections, both using standard link aggregation control

protocol (LACP). Blade Networks offers HotLinks, an active/passive redundancy mechanism for inter-switch links that works with or without spanning tree.

Cisco's Nexus switches offer Virtual Port Channel (VPC), which allows definition of logical server and inter-switch links that encompass multiple physical ports. Extreme's Summit x650 supports standards-based Ethernet Automatic Protection Switching (EAPS) for active/passive failover between switches, while the HP ProCurve switch uses a proprietary active/active mesh protocol to balance loads with upstream switches.

Arista's MLAG offers the twin benefits of being standards-based (it uses LACP and thus should work with any other device that supports link aggregation) and operating in active/active mode. Of the proprietary mechanisms, Cisco's VPC also offers interesting potential for carving out multiple virtual switching and routing domains within or across multiple Nexus switches.

To read the rest of David Newman's detailed review, click here: Storage/data convergence: Cisco stands alone

Network World *gratefully acknowledges the test equipment vendors that supported this project. Spirent Communications supplied its Spirent TestCenter Hypermetrics 10 gigabit Ethernet test modules for this project, along with extensive engineering and logistical support from Mark Hall, Gary Hansen, Mike Kanada, Michael Lynge, Himesh Mehta, Paul Mooney, Timmons C. Player, Howard Turner and Jurrie Van Den Breekel. Thanks too to Fluke, which provided Fluke 322 and 335 clamp meters for measuring power consumption.*

Newman is a member of the Network World Lab Alliance and president of Network Test, an independent test lab and engineering services consultancy. He can be reached at dnewman@networktest.com.

NW Lab Alliance

Newman is also a member of the Network World Lab Alliance, a cooperative of the premier reviewers in the network industry each bringing to bear years of practical experience on every review. For more Lab Alliance information, including what it takes to become a member, go to www.networkworld.com/alliance.

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

> **Must read: 11 hidden tips and tweaks for Windows 10**

View Comments

YOU MIGHT LIKE

Uber might charge you extra for making your driver wait

The Endless Mini $79 PC: A desktop for the masses

Storage/data convergence: Cisco stands alone

Most vendors have at least some virtualization features, including support for virtual switches

By David Newman | Follow
Network World | Jan 18, 2010 12:00 AM PT

Data center switches offer the promise of combining previously separate data and storage networks, thus reducing capital expense costs. However, the switches tested take different approaches to convergence.

Until the introduction of 10G Ethernet switches, data centers usually contained at least two separate networks: Fibre Channel for storage and Ethernet for data traffic. Faster Ethernet switches make it possible to combine traffic, perhaps by encapsulating storage traffic as FCoE or other means. But just shoving encapsulated Fibre Channel through an Ethernet port may not work well.

Fibre Channel and FCoE have stringent loss and delay requirements. Both use a credit-based system in which loss should never occur. Contrast that with Ethernet, which is stateless and doesn't attempt to retransmit lost frames. Further, Fibre Channel and FCoE require deterministic (predictable) latency and jitter; Ethernet doesn't.

To accommodate storage traffic on Ethernet, the IEEE has developed new mechanisms that aim to deliver Fibre Channel-like service levels over Ethernet. The most important of these is priority-based flow control (PFC), described in IEEE 802.1bb and supported by the Blade, Cisco and Dell switches. PFC extends Ethernet flow control to work on a per-VLAN priority basis instead of a per-port basis. One or more of those priorities can carry storage traffic.

A related mechanism is 802.1Qau congestion notification, which tells a given traffic class's transmitters to back off as queues fill. And yet another specification is DCBX, the data center bridging extensions to the logical link discovery protocol (LLDP), supported in the Blade and Cisco switches for exchanging capabilities information such as PFC support.

We didn't assess Fibre Channel or FCoE performance in this test because only Cisco's Nexus 5010 fully supports these protocols. If native support of Fibre Channel and/or FCoE are key requirements, the Nexus is the only choice in this test. (For the record, Brocade also makes a top-of-rack switch with full FCoE support, but the vendor declined to submit it for testing.)

Virtualization support

Virtualization is probably the biggest differentiator between data center and plain-vanilla Ethernet switches. Most vendors have at least some virtualization features, including support for virtual switches.

Arista's vEOS is a virtual extension of the vendor's Linux-based operating system that ties into VMware's Vsphere management platform. It allows any switch, physical or virtual, to be managed using the same image with the same capabilities.

The Blade, Cisco and Extreme switches all keep track of attached virtual machines and migrate associated policies – such as VLAN IDs and access control lists – as the VM moves to a new host.

Blade and Cisco offer a virtual network interface card (NIC) capability, though they work in different ways. Blade's approach divides a single 10G server Ethernet connection into four virtual channels that appear as separate NICs.

Cisco Nexus switches use virtual NICs together with VN-Link and VNtag technologies to map VM policies to VMs. For Fibre Channel and FCoE connections, the Cisco Nexus switches also support mechanisms called N-port virtualization and N-port identifier virtualization to reduce the number of Fibre Channel domain IDs and map multiple IDs to a single port.

See next part: Management and usability: Extreme goes its own way

Return to main test.

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

Uber might charge you extra for making your driver wait

The Endless Mini $79 PC: A desktop for the masses

Fixing Windows 10's privacy problems

NextAdvisor

NETWORKWORLD

Sign In | Register

NETWORKWORLD

Management and usability: Extreme goes its own way

Power consumption: Blade is most efficient

By David Newman Follow

Network World | Jan 18, 2010 12:00 AM PT

For most switches tested, the main question when it comes to usability is "Can you speak IOS?" Every switch tested except Extreme's has a command-line interface (CLI) that resembles Cisco's. That's a smart design choice considering more network engineers are conversant in IOS than any other environment.

Surprisingly, we found the CLI in the Cisco Nexus 5010, which looks like IOS but is actually a process running on the switch's NX OS, to be less convenient in some ways than IOS itself.

For example, IOS uses one simple command to download a configuration from a TFTP server. NX OS issues a warning that the command has been deprecated, and instructs the user to execute four other commands instead. Similarly, enabling jumbo frame support requires numerous policy-map and service-policy commands, compared with a single command in some Cisco Catalyst switches. Cisco says it's fixing the TFTP issue. We're glad to hear that; this is networking made cumbersome.

On the plus side, the Linux-based NX OS is modular, so a problem with any particular process won't bring down the entire system. That's a big improvement over monolithic IOS, and should enhance stability.

Arista's EOS also runs on Linux, and does more than any other switch tested to make Linux features available to users. This isn't just piping command output to any Linux program, handy as that is. The command set also allows network managers to drop into a Bash shell and run virtually any Linux command – including applying bug fixes without a reboot, a unique feature in this test.

Extreme's XOS is the only CLI tested not to emulate Cisco's IOS. We found XOS to be relatively simple to learn and use. That said, there were some fit-and-finish issues; for example, the switch allows swapping between primary and secondary configuration files but would not load a configuration using a different filename, even though the documentation says this is possible.

All switches tested support standard SNMP management information bases (MIB), and all except Arista's support the remote monitoring (RMON) MIB. All switches except Arista's also support at least some management methods over pure IPv6 networks. The Arista, Cisco, Extreme and HP switches also allow hierarchies of administrators to be defined with different privilege levels.

Power consumption: Blade is most efficient

Given that electricity is a major cost for any data center, it's not surprising that networking vendors have started to focus on power efficiency.

We measured power draw in this test using four configurations: With 12 ports idle and fully loaded, and again with 24 ports idle and fully loaded. We ran the 12-port tests to determine if switches saved additional power when only some switch ports have fiber-optic transceivers in place, as often happens in initial deployments. (We used 10GBase-SR transceivers in all tests.)

To measure maximum draw, we configured a Spirent TestCenter traffic generator/analyzer to offer 64-byte frames at line rate in a fully meshed pattern. That's a highly stressful pattern, and should produce power-draw readings at or near the maximum possible for any given switch. We then used a Fluke True-RMS 335 clamp meter to measure power consumption.

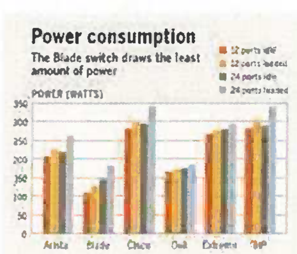

The clear leader in the power-measurement tests was Blade's G8124. It drew less power in the worst case (181 watts with 24 fully loaded ports) than most other switches' best-case numbers (all other switches except Dell's drew more power in the 12-idle-ports tests). Notably, the G8124 was the only switch to use proportionately less power with fewer transceivers attached.

The Dell and then Arista switches were the next most efficient across all test cases. The Cisco and HP switches were the least efficient of the bunch.

MAC address capacity: HP ProCurve sets the bar

MAC address capacity is a key differentiator between data center and enterprise switches. It's long been an article of faith in enterprise network design that broadcast domains – and thus the size of switches' MAC address tables – should be kept small through the use of routers and subnetting.

Virtualization changes all that. With migration tools such as VMware's Vmotion, all virtual machines must reside within the same broadcast domain. After all, a virtual machine in a given IP subnet won't be visible if it moves into a different subnet.

As a result, data center switches need large MAC address tables. In the old enterprise model, a top-of-rack switch might learn around 40 MAC addresses per server rack, plus a few MAC addresses for attached routers.

In the new data center, each server in that rack might accommodate up to 32 virtual machines (likely to rise to 50 or more using technologies such as Cisco UCS). That's 1,280 MAC addresses per rack – and in a large data center, there may be dozens or hundreds of such racks. Considering that all virtual machines need to see all others in Vmotion environments, switches may need to learn tens of thousands of addresses.

In our tests, the clear leader in MAC capacity was the HP ProCurve switch, which forwarded traffic to nearly 65,000 unique addresses without flooding, more than double the capacity of the Dell and Extreme switches, which came next. The Blade, Arista and Cisco switches all learned roughly 14,000 to 15,000 addresses.

For data center applications not involving virtualizations, any of these capacity numbers is more than adequate.

See next part: Latency and jitter: Cut-through design pays off for Arista, Blade

Return to main test

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

Must read: 11 hidden tips and tweaks for Windows 10

View Comments

Uber might charge you extra for making your driver wait

The Endless Mini $79 PC: A desktop for the masses

Fixing Windows 10's privacy problems

NETWORKWORLD

Sign In | Register

NETWORKWORLD

Latency and jitter: Cut-through design pays off for Arista, Blade

By David Newman | Follow
Network World | Jan 18, 2010 12:00 AM PT

In many data centers, latency and jitter are the most important metrics. Even the small amounts of delay (latency) or delay variation (jitter) introduced by a switch can have a profound impact on application performance. This is different than with general enterprise switching and routing, in which speed-of-light propagation times dwarf the latency added by any given switch or router.

In contrast, devices in the data center are only a few meters apart, or less, so every microsecond counts. Also, there's often a simple business driver involved: The more transactions an organization can process in a given unit of time, the more revenue it can expect to realize.

To help reduce latency and jitter, some switches (Arista, Blade, and Cisco) use so-called cut-through switching. A cut-through device begins forwarding a frame after examining only the first 12 bytes of its Ethernet header. In contrast, a store-and-forward switch caches the entire frame before making a forwarding decision.

The Dell, Extreme and HP switches all used store-and-forward mode in these tests. Extreme's Summit x650 can be configured in either mode; Extreme's engineers opted to use store-and-forward switching in order to achieve lower latencies.

Cut-through designs typically deliver lower latency, but there are tradeoffs. The biggest issue is that cut-through switches will forward corrupted frames, since they don't wait to see if the checksum at the end of each frame is valid. A router or other store-and-forward device will keep corrupted frames from leaving the data center, but such traffic could be a problem inside the data center, especially in large broadcast domains.

Another possible concern is that relatively low latency often means relatively small buffers. This isn't a problem when moving traffic between pairs of ports operating at the same speed, but speed mismatches between ports (say, gigabit and 10G Ethernet) or congestion from many-to-one traffic patterns could cause frame loss earlier than with store-and-forward devices.

Some vendor data sheets claim lower latencies than those we measured. Those claims may be valid, but they're not necessarily the most meaningful numbers for end-users (see related story "Lies, Damned Lies and Latency").

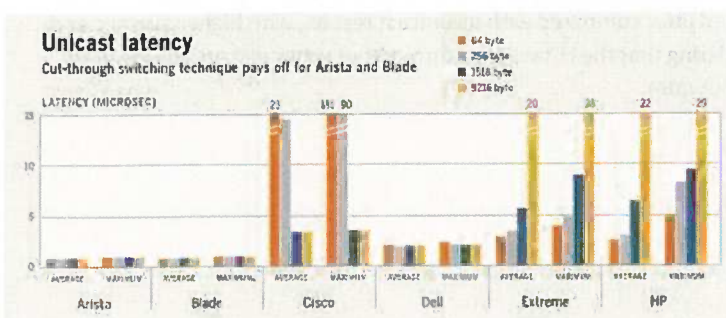

A cut-through design clearly paid benefits for the Arista and Blade switches, which delivered far lower latency across all frame sizes than their competitors. Blade's G8124 wins bragging rights with the lowest unicast latency – 750 nanoseconds with 64-byte unicast frames – but both the Arista and Blade devices consistently posted numbers around 800 nanoseconds in other tests.

Cut-through doesn't automatically translate to low latency, as the numbers from Cisco's Nexus 5010 make amply clear. When handling small frames, the Cisco switch delivered average latency that was 20 or more times higher than some other switches' maximum delays. Moreover, its maximum latency with 64-byte frames was a staggering 181 microseconds. At line rate, that means nearly 300 frames were in flight. That would be a high delay in a gigabit Ethernet switch, let alone a 10G Ethernet device intended for data center service.

Cisco attributes the latency numbers to superframing, a technique the Nexus 5010 uses to aggregate many frames into a single large unit for switching across the unit's crossbar component. The result, Cisco says, is a design that handles both Ethernet and storage traffic (in the form of FCoE), while accommodating FCoE's need for deterministic delay (that is, delay that's predictable and with little variation, or jitter).

Certainly it's true that the Nexus 5010's large-frame latency is far lower, around 3.37 microseconds, and that jitter is also much smaller. With large frames, the Cisco switch's maximum latency was 3.45 microseconds, only 80 nanosec higher than the average. This validates Cisco's claim of delivering deterministic delay and jitter.

On the other hand, the best-case latency for the Nexus 5010 is still more than four times higher than that of the Arista or Blade switches – and maximum delays for both those switches also were around 80 nanoseconds higher than their averages too, regardless of frame size. With the Cisco switch, in contrast, superframing only limits jitter with large frames. There's a bit of apples-and-oranges in this comparison; the Nexus switch handles Fibre Channel and FCoE in addition to Ethernet, and the other switches don't. Still, the simpler Ethernet devices clearly delivered lower latency and jitter across all frame sizes.

The Extreme and HP switches both exhibited significantly higher latency with large unicast frames, especially jumbo frames. Curiously, this was not the case with Dell's PowerConnect switch, even though it uses the same store-and-forward technique as the Extreme and HP devices. Dell's switch delivered very predictable latency for unicast traffic across all frame sizes, with average delays of less than 2 microsec and maximum delays around 200 nanosec higher.

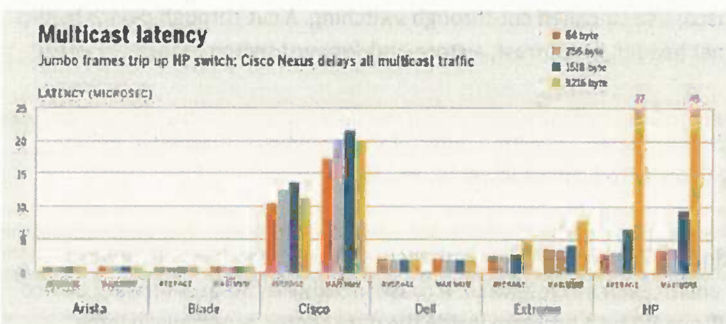

Multicast latency and jitter was virtually identical to unicast for the Arista, Blade and Dell switches. The Arista and Blade switches both achieved the lowest latencies we saw in the entire test, around 740 nanosec when handling 64-byte frames. Dell's numbers were also very consistent across all multicast frame sizes, and very similar to its unicast results.

Multicast latency for Cisco's Nexus 5010, while certainly lower than in the unicast tests, was still high vis-à-vis most other switches. Apparently superframing does not play a major role in multicast switching, since we observed relatively flat average and maximum delays across all frame sizes.

HP's ProCurve switch also turned in very different multicast latency and jitter compared with its unicast results, with higher average and maximum delays for jumbo multicast frames. This is surprising considering that the HP switch's throughput was substantially lower for multicast than for unicast traffic (latency is measured at the throughput rate).

See next part: **Link aggregation: Arista, Blade and Cisco fare best**

Return to main test.

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

> Must read: 11 hidden tips and tweaks for Windows 10

View Comments

Uber might charge you extra for making your driver wait

The Endless Mini $79 PC: A desktop for the masses

Fixing Windows 10's privacy problems

NETWORKWORLD

Sign In | Register

NETWORKWORLD

Link aggregation: Arista, Blade and Cisco fare best

By David Newman | Follow

Network World | Jan 18, 2010 12:00 AM PT

At least until the standards are done for 40G and 100G Ethernet, link aggregation will remain the preferred high-speed uplink method. Link aggregation, described in IEEE standard 802.3ad, binds multiple physical ports into a single, larger logical port – which in turns makes it easier to provision and manage high-speed flows.

Although link aggregation combines ports, it's not necessarily the case that throughput will scale linearly with port count. Fairness, or the uniformity of flow distribution across members of a link aggregation group (LAG), is always a concern.

Typically, switches inspect flow contents and use a hashing algorithm to assign each flow to a given member of a LAG. For example, every flow entering an eight-port LAG can take one of eight possible paths. A simple 3-bit hashing algorithm might inspect frame contents to determine where to assign the frame, because a 3-bit hash has eight possible outcomes. For example, the algorithm could inspect just the last 3 bits of each frame's source and destination MAC addresses to make a hashing decision.

To test the fairness of flow distribution, we configured each switch to support an eight-member LAG, and set up Spirent TestCenter to offer seven ports worth of unidirectional traffic across the LAG. The Spirent TestCenter test instrument emulated a total of 1,050 hosts, each using pseudorandom MAC addresses.

We counted the frames forwarded by each LAG member, and calculated the standard deviation of frame counts across all LAG members. We then administratively shut down one of the LAG members and again offered the same seven ports' worth of traffic across the remaining LAG members. While both the eight- and seven-port test cases measured fairness of distribution across LAG members, this second test also examined how each switch redistributes flows when LAG membership changes.

In both tests, we offered traffic at just 10% of line rate. Why? Because LAGs always introduce some unfairness of flows across LAG members, and the point of these tests is to find out how much. In situations with big disparities in frame counts between LAG members, frame loss can result if the offered load is high enough. Given the low offered load, we were confident that any disparities in frame counts were solely due to hashing across LAG members.

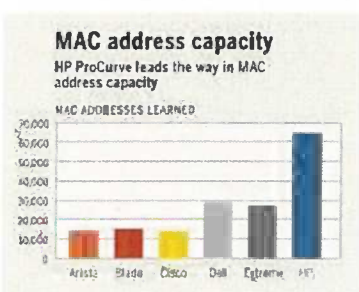

The results showed sizable differences between products. In the eight-port LAG case, Cisco's Nexus 5010 had the lowest standard deviation, meaning the most uniform distribution of frames across LAG members. The Dell and Arista switches were next most efficient.

In the seven-member case, the Cisco switch was about three times less uniform than in the eight-member case. And frame distributions were even more uneven for the Dell and Extreme switches, with standard deviations 11 and six times larger, respectively, than in the eight-member test case.

With the Dell and Extreme switches, if a LAG member goes away, all its flows simply migrate to one other LAG member. That led to a big jump in frame counts on that single member.

The most efficient switches in the seven-member LAG tests by far were Blade's and Arista's. Both actually improved hashing efficiency, something none of the other switches did. Arista's switch showed the least variation between eight- and seven-port cases, while the Blade switch showed the greatest improvement in the seven-port case.

See next part: Multicast group capacity: Extreme comes out on top

Return to main test.

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

Must read: 11 hidden tips and tweaks for Windows 10

View Comments

Is the Elio P5 the car of the future?

Why can't the FBI unlock that iPhone?

Zosia Mamet & Kat Dennings in Missadventure's Season 2 - Watch

NETWORKWORLD

Multicast group capacity: Extreme comes out on top

Multicast join/leave delay: Arista and Dell are swell

By David Newman | Follow
Network World | Jan 18, 2010 12:00 AM PT

As data center managers consolidate and virtualize their servers, the next order of business becomes moving all that traffic. Enter top-of-rack data center switches that offer speed, scalability, redundancy, virtualization support and other features not available in garden-variety Ethernet switches.

Because many data center applications today make use of IP multicast, scalability is naturally a concern. For layer-2 switched environments, the main measure of IP multicast scalability is group capacity, or the number of Internet group membership protocol (IGMP) snooping entries a switch can keep track of.

Switches snoop or listen for IGMP group membership reports and then should forward traffic only to those ports where hosts have subscribed to specific groups.

Extreme's Summit x650 was the clear leader in multicast group capacity, successfully forwarding traffic to 6,000 groups. The Arista and HP switches were next, each forwarding traffic to 2,047 groups. Cisco's Nexus was a little behind them with 2,000 groups, while the Blade and Dell switches each supported 1,024 groups.

Multicast join/leave delay: Arista and Dell are swell

There can be real money tied to the speed at which a switch processes requests to join or leave multicast groups. Financial services companies depend on getting quotes quickly; search engine clusters need to keep synchronized for their companies to realize ad revenue; and ISPs with IPTV offerings use multicast to add and remove subscribers. For all these applications and many more, the speed of multicast group join and leave processing is a key consideration.

As with the multicast throughput and latency tests, we used 989 groups and had hosts on all receive ports join all groups. This time, however, we set aside one port as a monitor port that should neither transmit nor receive multicast traffic. This extra port was a check against flooding.

Figure 5: Multicast join/leave delay

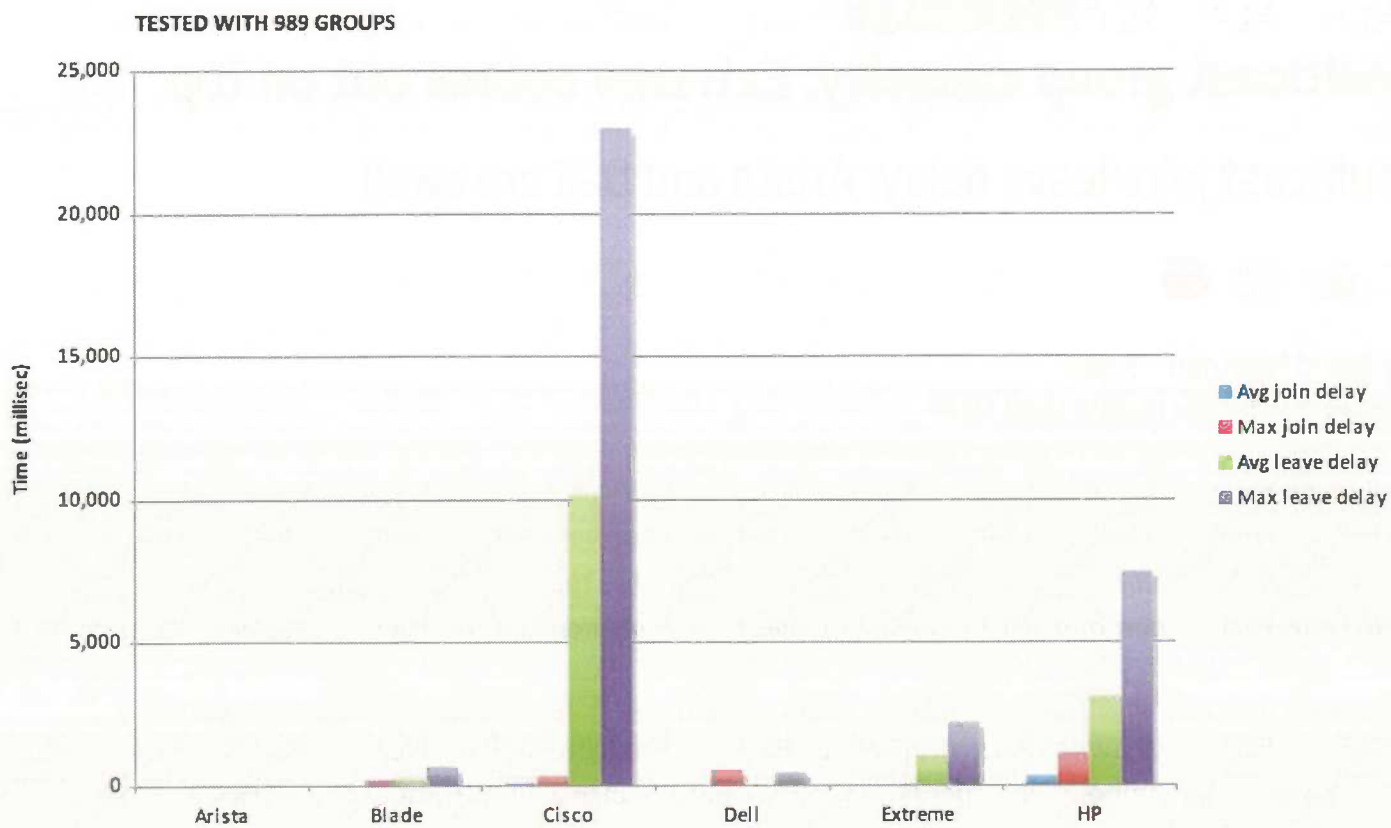

Dell's PowerConnect 8024F had the lowest average join and leave times, followed closely by the Arista and Blade switches. However, Arista's 7124S was more consistent across the board, with the least variation between average and maximum join and leave times. This is largely a function of control-plane processing power, and reflects Arista's use of a dual-core 1.8-GHz x86 CPU, a powerful processor for a top-of-rack switch.

The HP and Cisco switches had much higher join and leave delays than other devices. In the worst case, Cisco's switch took nearly 23 seconds to process a leave message for one group and averaged about 10 seconds to process each leave message. The Nexus join times were much lower, in the low hundreds of milliseconds.

Even if one assumes that leave delays are less important than join delays, there's also another issue with the Nexus switch: It leaks multicast traffic when its control plane is busy. We noticed that whenever the switch was processing responses to IGMP queries, it flooded multicast traffic to that last port – the one with no multicast subscribers, the one that never should have received multicast traffic.

Cisco says this is working-as-designed behavior. It appears that when the switch's CPU is very busy (as it is when processing responses to IGMP queries), the switch will flood multicast traffic to all ports, even those with no subscribers attached. This may be intentional, but at the same time we saw no flooding when running the same test on two Cisco Catalyst switches.

We also wondered if the multicast leakage might be just an artifact of our stress testing, but that turned out not to be the case. We still observed flooding even when we cut down the number of multicast receive ports from 22 to just one and reduced the traffic rate from line rate to 10% of line rate.

After testing concluded, Cisco reproduced the leakage issue and says it's working on a fix.

Forward pressure: No cheaters allowed

The final set of tests looked at an area where lots of switch makers used to cheat: Forward pressure, or the practice of transmitting frames spaced too closely together. This was a big problem in the early days of Ethernet switching, when makers of half-duplex switches used to try to win collision battles by putting a frame on the wire sooner than the attached station could.

While we're long past the days of half-duplex switching, it's still important for switches not to run too fast or too slow. A too-fast switch can cause an attached host, switch or router to drop traffic. A switch that transmits too slowly may itself drop frames over time, given a buildup of traffic from a faster link partner.

The IEEE 802.3 Ethernet specification allows for some leeway in clock rates. It says every Ethernet interface must tolerate variations in clock speeds of 100 parts per million (ppm). In traffic terms, that translates to about 1,488 fps when 64-byte frames are involved. Expressed another way, 10G Ethernet devices should run at 14,880,952 fps with 64-byte frames, but it's OK if the rates are up to 1,488 fps faster or slower than that rate.

To find each switch's maximum transmission rate, we configured Spirent TestCenter to cheat. As a test instrument, it can deliberately transmit traffic with illegally small gaps between frames. Nominally, Ethernet is supposed to operate with a 12-byte gap between frames; for this test, we offered traffic with 11- and then 10-byte gaps to determine maximum forwarding rates.

All switches passed this test, meaning none forwarded traffic above the 100-ppm limit. At the same time, we observed a wide range forwarding rates within the 100-ppm tolerance. As noted in the throughput discussion, the Cisco Nexus 5010 forwarded traffic a bit slower, around 17 ppm, below line rate. The Dell switch was very close to nominal line rate, at around 2 ppm over line rate. Next fastest was the Blade switch at around 23 ppm above line rate. The Extreme and Arista switches ran the "hottest," at around 45 and 50 ppm over line rate respectively.

Return to main test

David Newman

Newman, a Network World Test Alliance partner, isÂ president of Network TestÂ in West Lake Village, Calif. He can be reached at

Must read: 11 hidden tips and tweaks for Windows 10

View Comments

YouTube's Content ID program finally provides for ad revenue during disputes

Is the Elio P5 the car of the future?

Why can't the FBI unlock that iPhone?